

Active neural representations: Neurophysiological data and conceptual implications

Marius Usher

Department of Psychology, University of Kent at Canterbury
Kent, CP2 7NP, UK
e-mail: M.Usher@ukc.ac.uk

Ernst Niebur

Krieger Mind/Brain Institute and Dept. of Neuroscience
The Johns Hopkins University, Baltimore, MD 21218, USA
e-mail: niebur@jhu.edu

April 12, 1997

Abstract

We discuss the nature of cognitive representations and present a scheme for the encoding of information which accounts for both categorical and graded aspects of cognitive events. Accordingly, object categories are mapped onto the identity of cell populations via the process of categorical perception, while graded aspects such as vividness and confidence level are mapped onto the rate and degree of synchrony of neural populations. We show that the process of categorical perception can solve the problem of reference of perceptual-cognitive states, and we demonstrate by computer simulation models that synchrony can encode for grouping and binding of coherent information. In these models, recurrent connections play a major role in generating synchronised activity in cell populations which receive consistent sensory stimulation, and in generating temporal fluctuations in their discharge rate. We suggest that such fluctuations underlie autonomous computations reflecting fluctuations in the certainty (confidence level) of perceptual and cognitive hypothesis.

Introduction

Despite the fact that mental representations are a central concept in cognitive psychology and in neuroscience, a basic understanding of their nature and function seems to be lacking. It is possible that this state of affairs is responsible for the fact that, within the philosophy of mind, the concept of mental representation has been challenged as being inconsistent. Following Wittgenstein (1958) and Putnam (1981), some philoso-

phers argued that there is no way by which mental or brain traces can satisfy the problem of reference (or intentionality), namely, justifying how brain traces refer to (i.e., are about) objects and events in the external world. We will also discuss the nature of the neural code, a subject of intense controversy: does the brain encode information in the firing rates of cell-populations or in the exact timing of neural discharge patterns (Softky, 1995; Shadlen & Newsome, 1994, 1995), or in yet another way? Evidently, recent neurophysiological data indicates that the classical scheme based on a bottom-up hierarchy of encoding cells is inadequate; neural representations seem to be distributed and most of the input is contributed by feedback and lateral connections rather than by afferent connections.

The scope of this report is twofold. First, on the basis of neurophysiological data, we propose a scheme for the "neural code". A main characteristic of this scheme is that neural representations are not passive entities activated by external input, but are intrinsically related to active and autonomous brain processes, mediated by feedback and lateral connections. Second, we will try to demonstrate that this neural scheme can meet the challenge of the reference problem, providing thus a basis for a representational theory of mind.

Our paper is organized as follows. In the next section we review some conceptual arguments that are typically raised against the representational theory of mind and we try to show that active brain processes, such as *categorical perception*, can overcome those problems. In the third section we

introduce a scheme for neural coding which accounts for both discrete categories and graded (intensity like) characteristics of cognitive and perceptual representations. Finally, in the fourth section, we further develop this scheme by using computational models and show how temporal aspects of neural patterns, such as synchronisation of discharge, might contribute to more subtle perceptual processes, such as grouping and binding of coherent information.

Mental Representations And Categorical Perception

The problem of mental representation

Simply put, the problem is to explain how can brain states (or any other object) represent (refer or point to) objects in the world (other than by convention¹). There are in fact only two principles that have been seriously considered by philosophers to account for a naturalistic and reductive explanation of mental representation: *resemblance* and *causation* (Fodor, 1990; Crane, 1995). However, theories based on both of these principles, have been shown to suffer from severe difficulties (see Crane, 1995, for a review of such theories).

The problem with the resemblance principle can be illustrated with the following example used by Putnam (1981). An ant moves across the sand leaving a trace which by pure accident resembles the shape of Winston Churchill. It seems quite compelling that despite its similarity this accidental trace does not represent the person of Winston Churchill, in itself. Thus resemblance cannot be a sufficient condition for representation².

The causation principle seems a much better candidate for a theory of representation. Accordingly, X represents Y if and only if X is caused by Y³. This seems indeed to be the missing ingredi-

ent in the previous example: the trace of the ant has no causal relationship to the person of Winston Churchill, while a perceptual brain state is caused by the presence of its perceptual object. However, a theory of representation based on causation alone is still insufficient, as it suffers from two ubiquitous problems: the problem of error and the problem of interpretation.

To illustrate the former, notice that a valid theory should be able to account for errors in the representation (e.g., under short exposure conditions, one might miss-perceive the letter "A" for an "H"). Even in such a case, however, the underlying brain state is causally linked to the miss-perceived object. It is thus not possible to account on the basis of causation alone for the difference between veridical and erroneous perception. Consider also the perceptual process taking place when perceiving an ambiguous figure. Since mental representations can not rely on external interpretation (to avoid infinite regression), the perceptual process itself should provide the interpretation. Definitely there is nothing in virtue of which a merely causal process, such as a photographic projection, could provide an interpretation of the ambiguity. Thus we see that in order to refer (or represent) a special interpretation process is required. Important questions arise: What are the essential characteristics of such an interpretation process? Is there a specific process of interpretation taking place in the brain during the perception of objects, different from a photographic projection and which could account for the ability of mental/brain states to represent external objects?

We believe that indeed, brain states can represent external objects, not due to some magical brain properties, but because of special interpretation-like causal but *stochastic* processes (for an account of probabilistic causality see, e.g., Eells, 1991), taking place during sensory perception. Moreover, we believe that such a process is now well understood in cognitive psychology and neuroscience. We turn now to the process of *categorisation* and *categorical perception* which has been demonstrated both in human and in animal subjects (Harnad, 1987; LaMotte & Mountcastle, 1979; Wytttenbach, May, & Hoy, 1996), and which we believe to have the necessary characteristics for solving the intentionality problem.

¹While convention is a reasonable explanation for many kind of representations, say, words, such a solution is totally unsatisfactory for brain states (which unlike words should be an instance of "underived" representation), due to infinite regression.

²One might also consider the fact that the words "Winston Churchill" spoken by a human being who is familiar with modern history do represent the person of Churchill despite of bearing no similarity to him. This could be taken to imply that "resemblance" is not a necessary condition, either; note however that word-representations are derived (i.e., they rely on convention).

³A technical detail used sometimes, is to require causation based on natural regularity (or "reliable indication") (Crane, 1995). As this will make no difference in the following discussion we use only the simple causation label

instead.

Categorisation and categorical perception

There are several ways in which the process of categorical perception is fundamentally different from a mapping such as a shadow or a photograph. Unlike those passive mappings which transform a continuity of objects to a similar continuity of projections (e.g., infinitesimal rotations or distortions of the objects are mapped into the projection), the process of categorical perception involves a discontinuous decision-like process: continuous families of similar stimuli are mapped (probabilistically, depending on environmental or perceptual noise) onto discrete categorical concepts. Thus unlike a photographic projection, categorical perception is an active interpretation process in which small irrelevant transformations (distortions) are ignored and which involves perceptual boundaries⁴. Computational models for perception and memory (Anderson, Silverstein, Ritz, & Jones, 1977; Hopfield, 1984) have also demonstrated how systems of neural elements can generate categorical perception, generating computations which are error-corrective due to *attractor* dynamics. As a consequence, perceptual states do not need an external interpretation; they are part of the interpretation.

To understand how the process of categorisation solves the problems of perceptual mental representations⁵, we should consider the *ensemble* of objects (characterised by some statistical distribution) that an organism with a cognitive system encounters and interacts with during its life. One can then propose that a brain categorisation state represents the most *typical* object that is likely to be mapped to (categorised by) that state (this is well defined by the statistical characteristics of the ensemble of objects and by the stochastic categorisation process⁶). This scheme bears some similarity to the approach of Dretske (1981) who emphasised the role of the learning period for providing the reference of representational states; the statistical ensemble reflects the learning experience

⁴Of course, some simple forms of “categorical perception” takes place even in simple physical systems such as a thermostat which maps all temperature conditions above a specified boundary to a “hot” category and the temperatures below the boundary to a “cold” category.

⁵It is beyond the scope of this paper to provide an account of cognitive representations beyond the perceptual domain. One should note however, that perceptual representations could provide a basis for building a full repertoire of cognitive representations using a hierarchical scheme (Devitt & Sterelny, 1987).

⁶Such a “typical object” can correspond to what cognitive psychologists call a *prototype*. One should note however, that a prototype could depend on context, unlike in naive categorisation theories.

of the organism. However, unlike in that scheme (Dretske, 1981), mis-representations are naturally accounted for without the need to rely on the assumption that the learning stage is errorless, in order to distinguish true from incorrect reference. They are the instances when an object has been categorised (due to noise in the environment or in the perceptual system) to a categorisation state which is atypical for it. Such occurrence should be rare (otherwise the categorisation state would be a typical one)⁷ The process of categorisation can thus be seen as an instance of statistical hypothesis testing: finding which object (out of a set of alternatives that depend on context) is the more likely to have caused the perceptual state.

There are thus important differences between the ant’s trace in the sand and a brain trace which is the result of categorical perception. First, unlike the trace in the sand, the perceptual state is causally related to the perceived object. Second, unlike the trace in the sand, the perceptual state reflects an interpretation of the external reality: classifying families of sensory stimuli according to perceptual boundaries which reflect “typical” objects in the external world. Given the activation of such a classification-state one can predict that there is a relatively high probability that it was actually generated by the represented object. Surely, such perceptual predictions might be wrong: the brain state might be generated as result of an optical illusion, a drug induced hallucination, etc. Even under such “atypical” conditions, however, the perceptual brain state refers and represents the object which would generate the trace in typical situations and which is predicted by the process of perception.

Although Putnam’s example, illustrated above, dismisses similarity as a necessary relation for the faculty of representation, the use of similarity in the perceptual categorisation process can improve the efficiency of the representation. This is due to the fact that the cognitive/perceptual system needs not only to represent objects in isolation but also to represent essential relationships between the objects. A perceptual scheme where traces bear some resemblance to the perceived object will automatically conserve similarity relationships among objects. In addition, the simi-

⁷It is interesting to note that this mechanism for causality satisfies the principle described by Fodor (1990) as asymmetric dependency. If object A is misrepresented by the representation of object B, the causal process involved is mediated (thus dependent) on the representation of B (its attractor pattern), while the opposite is not true (categorizing object A as A does not depend on the representation of object B).

larity of an actual object to some prototype will generally determine the probability that the corresponding object will be mapped to the representation of that prototype.

Categories and intensity

Until now we focussed on one property of cognitive state: their referent. This is however not the only aspect of mental states. Mental states are also essentially characterised by their phenomenology: how it *feels* being in that state. For example a pain can be intense or shallow, and the same verbal representation feels different if it reflects a belief, a hope or a desire. It is beyond the scope of this work to provide a complete account to phenomenology, however as a first step we will try to account to a single aspect of it: the *intensity* of the perceptual hypothesis (vividness, certainty, or confidence level); unlike categories which are discrete entities, intensities are analog variables. In the following section we will discuss several coding schemes which address both the reference and the intensity of cognitive states.

Coding Schemes

The discussion of irregularity of discharge in cortical neurons (Softky & Koch, 1993) has re-opened and further stimulated an old controversy (e. g., Werner & Mountcastle, 1963; Poggio & Viernstein, 1964; Bullock, 1970): Is the irregularity of discharge the manifestation of a complex and hidden signal, or does it reflect the noisy, stochastic nature of neural information processing? These two approaches have been contrasted by Ferster and Spruston (1995) as rate-code *versus* temporal code hypothesis.

In this context it is useful to distinguish between two issues which have often been confounded in the literature (Softky, 1995; Shadlen & Newsome, 1995; Ferster & Spruston, 1995): The first is whether temporal patterns (beyond firing rates) are being used in the coding process. The second, and independent, issue is the role that such temporal patterns play in the representation scheme.

The following possibilities can be considered:

1. **Classical population rate-code.** Categories (stimulus-identity) are represented by the identity of neural populations (labeled-line hypothesis), while analog intensity values map onto the rates of response (or onto the number of responding neurons) *within the same population*. Such a population rate code

is implicitly assumed in most of the connectionist models where the activation of units is typically assumed to represent the firing rate of neural populations (Wilson & Cowan, 1972; Anderson et al., 1977; Rumelhart, McClelland, & the PDP Research Group, 1986; Hopfield, 1984; Amit, 1989).

2. **Pure-temporal codes.** Categories are represented by different temporal patterns of activity *within the same population* (e.g., like in a Morse code (Softky, 1995)) and intensity is mapped onto the number of cells which signal this pattern; the identity of the cells is irrelevant.
3. **Hybrid rate-temporal population code.** Categories are represented by the identity of neural populations. Temporal patterns (rate or synchrony) signal the intensity level. In addition, synchrony of neural populations encodes *relational* properties (i.e., binding) between the elemental properties (features) represented by those populations (Gray, König, Engel, & Singer, 1989; Crick & Koch, 1990a, 1990b; Singer & Gray, 1995). This scheme can be thought of as a synthesis of the first two; from the first scheme, it inherits the hypothesis that elemental contents are represented by identity of cell populations; from the second, that temporal patterns play a role in the encoding.

As it seems, most discussions have contrasted the first two of these coding schemes. Since they did not consider the third alternative, they had the tendency to interpret any data suggesting a role for temporal patterns in information processing as support for the second coding-scheme: that the same cell (or cell population) can represent various objects (different contents), depending on its temporal pattern (Softky, 1995; Ferster & Spruston, 1995). We propose that the third coding scheme, the hybrid rate-temporal code is best supported by experimental data, and is likely to be computationally advantageous given the resources and limitations of physiology. Additional arguments and further evidence for this claim are presented in a separate article (Usher & Niebur, 1996).

Recurrent Connections And Temporal Patterns

Synchrony and binding

The population rate-temporal code is consistent with a role for temporal synchrony (on a time scale of 5-10 msec) in coding of relational aspects (binding; Gray et al., 1989; Roelfsema, König, Engel, Siretenau, & Singer, 1994). This scheme is supported by the wealth of experimental evidence which indicates that categorical properties of stimuli are represented by cell-populations. The use of synchrony and phase-relationships for coding relational properties can greatly improve the representational power of this scheme (Crick & Koch, 1990a, 1990b; Milner, 1974; von der Malsburg, 1981).

In our previous models (Usher & Niebur, 1995; Usher, Stemmler, & Niebur, 1996) we have also shown that synchronised neural activity corresponding to grouping of coherent visual information arises in simple models of the primary visual cortex due to recurrent lateral connections. For instance, Figure 1 illustrates an elementary form of binding in a model of visual cortex, implemented as a two-dimensional sheet of integrate-and-fire neurons with local and long-range connectivity. To demonstrate binding, we computed the cross-correlations between local field potentials (LFP) recorded (in the simulation) at two sites which are stimulated either by the same or by orthogonal orientation gratings (within the same distance between each pair of sites). As shown in Figure 1 (bottom), the correlation between cells receiving the same input is significantly stronger than that between cells that receive different inputs. One should notice that this effect is stimulus dependent and can not be explained by the simple fact that connected cells synchronise (Mirolo & Strogatz, 1990). In fact, the degree of correlation between two given cells depends on their stimulation and is not only a function of the connectivity.

Synchrony as an intensity code

Strong support for a role of synchrony in the representation of stimulus-strength is provided by recent reports by Decharms and Merzenich (1996), which are consistent with a hybrid rate-temporal code. Accordingly, a cell population can signal the presence of the stimulus which it encodes either by an increase in firing rate or by an increase in synchrony (within the population); under either condition, the presence of the stimulus can be transmitted to cells further up in the processing stream.





Figure 1: Activity in neuronal array. **Top:** Array of 100×100 cells where the gray-level of each pixel corresponds to the preferred orientation of the respective cell (lightest gray: horizontal, light gray: 45° , dark gray: 90° , black: 135°). White pixels represent spiking cells. **Bottom:** Cross-correlation functions of local field potentials. The dotted curve is between LFP recorded at sites stimulated with the same orientation grating and the full line is between LFP recorded at sites stimulated with orthogonal gratings.

As a consequence, synchrony is signalling the intensity and not the contents of the represented stimulus (which is determined by the identity of the neural population). This view also agrees with the hypothesis that selective attention is implemented by increased synchrony of discharge within a neural population (Niebur & Koch, 1994)⁸.

Rate Fluctuations

We propose that the irregularity of neural discharge is due to a large extent to self generated

⁸Other temporal codes are possible; see Niebur, Koch, and Rosin (1993) for an example.

fluctuations in the firing rate of neural populations. In another article (Usher & Niebur, 1996) we show that experimental data are consistent with the predictions of network models which generate their own rate fluctuations (Usher, Stemmler, Koch, & Olami, 1994; Usher, Stemmler, & Olami, 1995; Hansel & Sompolinsky, 1996). Such self generated fluctuations arise due to the feedback connections, in particular for center-surround lateral connectivity, enabling the networks to respond fast by rapidly amplifying fluctuations in the external input (Usher et al., 1995). Such fluctuations might underlie internal computations and reflect fluctuations in the certainty and the confidence level of perceptual/cognitive hypotheses. Optical imaging techniques have recently revealed evidence for such fluctuations (Arieli, Sterkin, Grinvald, & Aertsen, 1996). In particular, Arieli et al. (1996) found that the variability of evoked visual responses is due to a large extent to the variability in the spontaneous activity, present previously to the sensory stimulation. In Figure , we show a very similar behavior in a simulation model based on (Usher et al., 1994, 1995) (see caption).

Discussion

The neural code proposed here makes a distinction between the encoding of contents and the encoding of intensity. According to our scheme, contents are encoded through the identity of neural populations. The perceptual processes using those populations are probably related to the phenomenon of categorical perception (Harnad, 1987), believed to play an essential role for the generation of mental representations (Harnad, Hanson, & Lubin, 1994; Harnad, 1994). This process relies to a large extent on recurrent connections, as illustrated in various computational models (Anderson et al., 1977; Hopfield, 1984). We have argued that this process provides the essential ingredient which enables cognitive representations within the cortex to refer to objects in the external world.

Neurophysiological data seems to support population codes: the category information is determined by the identity of the cell-populations and the fluctuating activity levels reflect fluctuations in the strength associated with the response (or the confidence level of a perceptual hypothesis), rather than different "contents". According to the proposed hybrid population code, fluctuations within a neural population can be interpreted as continuously changing intensities of perceptual/cognitive hypotheses. These fluctuations might reflect changes in confidence level or in the

certainty of the hypotheses, which occur even in the absence of a stimulus, within the spontaneous ongoing cortical activity (Arieli et al., 1996). This suggests an active view of the perceptual/cognitive system where even in the absence of inputs, various hypotheses are being continuously entertained at various degrees (and bound in various combinations), and information from the past is integrated within present responses, supporting an active view of cognition.

Acknowledgements

The work of EN is supported by the National Science Foundation and by the Markey Foundation.

Figure 2: (next page) Three instances of spontaneous (left column) and evoked (right column) activity. The evoked stimulus provides input to the upper half of the simulated cortical layer, and the evoked response is "measured" 20 msec after the stimulus input impinges on the layer. Note that a) responses are mainly concentrated around clusters of spontaneous activity (since spontaneous activity varies across trials this leads to variability in the evoked response, as reported by Arieli et al. (1996), and b) responses are stronger in the upper half-plane, functionally corresponding to an amplification of the input.

References

- Amit, D. J. (1989). *Modelling Brain Function*. Cambridge University Press, Cambridge.
- Anderson, J. A., Silverstein, J. W., Ritz, S. A., & Jones, R. S. (1977). Distinctive features, categorical perception and probability learning: Some applications of a neural model. *Psychological Review*, *84*, 413–451.
- Arieli, A., Sterkin, A., Grinvald, A., & Aertsen, A. (1996). Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science*, *273*, 1868–1871.
- Bullock, T. H. (1970). The reliability of neurons. *J. Gen. Physiol.*, *55*, 565–584.
- Crane, T. (1995). *The mechanical mind*. Penguin, London, UK.
- Crick, F., & Koch, C. (1990a). Some reflections on visual awareness. *Cold Spring Harbor Symp. Quant. Biol.*, *55*, 953–962.
- Crick, F., & Koch, C. (1990b). Towards a neurobiological theory of consciousness. *Seminars in the Neurosciences*, *2*, 263–275.
- Decharms, R. C., & Merzenich, M. M. (1996). Primary cortical representation of sounds by the coordination of action potential timing. *Nature*, *381*, 610–613.
- Devitt, M., & Sterelny, K. (1987). *Language and reality*. Blackwell, Oxford.
- Dretske, F. (1981). *Knowledge and the flow of information*. Blackwell, Oxford.
- Eells, E. (1991). *Probabilistic causality*. Cambridge University Press, Cambridge, UK.
- Ferster, D., & Spruston, N. (1995). Cracking the neuronal code. *Science*, *270*, 756–757.
- Fodor, J. (1990). *A theory of content and other essays*. MIT Press, Cambridge, MA.
- Gray, C., König, P., Engel, A., & Singer, W. (1989). Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature*, *338*, 334–337.
- Hansel, D., & Sompolinsky, H. (1996). Chaos and synchrony in a model of a hypercolumn in visual cortex. *Journal of Computational Neuroscience*, *3*(1), 7–34.
- Harnad, S. (1987). *Categorical perception: The groundwork of cognition*. Cambridge University Press, Cambridge, UK.
- Harnad, S. (1994). Does mind piggyback on robotic and symbolic capacity. In Morowitz, H., & Singer, J. L. (Eds.), *The mind, the brain, and complex adaptive systems*, pp. 203–220. Addison-Wesley Publishing Co, Reading, MA.
- Harnad, S., Hanson, S. J., & Lubin, J. (1994). Learned Categorical Perception in Neural Nets: Implications for Symbol Grounding. In Honavar, V., & Uhr, L. (Eds.), *Symbol Processors and Connectionist Network Models in Artificial Intelligence and Cognitive Modelling: Steps Toward Principled Integration*, pp. 191–206. Academic Press.
- Hopfield, J. J. (1984). Neurons with Graded Responses Have Collective Computational Properties Like Those of Two-State Neurons. *Proc. Nat. Acad. Sci., USA*, *81*.
- LaMotte, R. H., & Mountcastle, V. B. M. (1979). Disorders in somethesis following lesions of parietal lobe. *J. Neurophysiology*, *42*(2), 400–419.
- Milner, P. (1974). A model for visual shape recognition. *Psychol. Rev.*, *81*, 521–535.
- Mirollo, R., & Strogatz, S. (1990). Synchronization of pulse-coupled biological oscillators. *SIAM J. on Applied Mathematics*, *50*(6), 1645–1662.
- Niebur, E., & Koch, C. (1994). A model for the neuronal implementation of selective visual attention based on temporal correlation among neurons. *Journal of Computational Neuroscience*, *1*(1), 141–158.
- Niebur, E., Koch, C., & Rosin, C. (1993). An oscillation-based model for the neural basis of attention. *Vision Research*, *33*, 2789–2802.
- Poggio, G. F., & Viernstein, L. J. (1964). Time series analysis of impulse sequences of thalamic somatic sensory neurons. *J. Neurophysiology*, *27*, 517–545.
- Putnam, H. (1981). *Reason, truth, and history*. Cambridge University Press, New York.
- Roelfsema, P. R., König, P., Engel, A. K., Sireteanu, R., & Singer, W. (1994). Reduced synchronization in the visual cortex of cats with

- strabismic amblyopia. *European Journal of Neuroscience*, 6(11), 1645–1655.
- Rumelhart, D., McClelland, J., & the PDP Research Group (1986). *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, Vol. 1. MIT Press, Cambridge.
- Shadlen, M. N., & Newsome, W. T. (1994). Noise, neural codes and cortical organization. *Current Opinion in Neurobiology*, 4, 569–579.
- Shadlen, M. N., & Newsome, W. T. (1995). Is there a signal in the noise?. *Current Opinion in Neurobiology*, 5, 248–250.
- Singer, W., & Gray, C. M. (1995). Visual feature integration and the temporal correlation hypothesis. *Annu. Rev. Neurosci.*, 18, 555–586.
- Softky, W. (1995). Simple codes versus efficient codes. *Current Opinion in Neurobiology*, 5, 239–247.
- Softky, W., & Koch, C. (1993). The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J. Neurosci.*, 13(1), 334–350.
- Usher, M., & Niebur, E. (1995). Stimulus-dependent synchronization in a model of orientation-tuned spiking neurons. In Bower, J. M. (Ed.), *Computation and Neural Systems*, Vol. 3, pp. 341–346. Kluwer, Norwell, MA.
- Usher, M., & Niebur, E. (1996). Firing Patterns And Neural Codes: Models, data and implications. *in preparation*.
- Usher, M., Stemmler, M., Koch, C., & Olami, Z. (1994). Network Amplification of Local Fluctuations Causes High Spike Rate Variability, Fractal Firing Patterns and Oscillatory Local Field Potentials. *Neural Computation*, 6(5), 795–836.
- Usher, M., Stemmler, M., & Niebur, E. (1996). The role of lateral connections in visual cortex: dynamics and information processing. The UTCS Neural Networks Research Group, Austin, Texas. Electronic book, ISBN 0-9647060-0-8, <http://www.cs.utexas.edu/users/nn/web-pubs/htmlbook96>.
- Usher, M., Stemmler, M., & Olami, Z. (1995). Dynamic pattern formation leads to 1/f noise in neural populations. *Physical Review Letters*, 74(2), 326–329.
- von der Malsburg, C. (1981). The correlation theory of brain function. Tech. rep. 81-2, Max-Planck-Institute for Biophysical Chemistry, D-3400 Goettingen, Germany.
- Werner, G., & Mountcastle, V. B. (1963). The variability of central neural activity in a sensory system and its implications for the central reflection of sensory events. *J. Neurophysiol.*, 26, 958–977.
- Wilson, H., & Cowan, J. (1972). Excitatory and inhibitory interactions in localized populations of model neurons. *Biol. Cybern.*, 12, 1–24.
- Wittgenstein, L. (1958). *Philosophical Investigations* (Second edition). Blackwell, Oxford.
- Wyttanbach, R. A., May, M. L., & Hoy, R. R. (1996). Categorical perception of sound frequency by crickets. *Science*, 273, 1542–1544.